**Printing History**

| Date | Document | Release |
|---|---|---|
| 07/16/97 | 01-000000-100-0041 | Interceptor 1.0 Installation and User's Guide |

Printed in U.S.A.

# Table of Contents

# About This Guide

The WebSpective User's Guide is intended to help you quickly learn to make productive use of WebSpective functionality. In order to aid in this process, this guide uses a number of conventions that are exercised throughout the book. These conventions are described in this chapter, along with a list of tasks and topics with which a WebSpective User should be familiar.

## Chapter Topics

- Navigating the Manual
- Style Conventions
- Common Terms
- Feedback

## Navigating the Manual

In order to help you move more quickly through the manual, a table of contents and index are provided in addition to a topic listing that occurs at the beginning of each chapter. Topics in the topic listing appear in the order in which they occur in the chapter.

A numbering convention has also been instituted to help you find items in the book. In both the table of contents and the index, page numbers are preceeded by their chapter number to help you find them faster.

## Style Conventions

The following style conventions have been instituted to help you discern different kinds of information in the book:

`System text is used wherever a command-line entry is required`

**How-to headings**

How-to headings signify the beginning of a step-by-step operation. They are listed in italics in the table of contents to help differentiate them from regular sub-headings.

- **Bolded Items**—are items which are key to WebSpective's operation. Bolded items are usually configuration parameters or menu options.

## Common Terms

The following terms are used frequently throughout the guide. Make sure that you understand how these terms are used in this guide as their definitions may differ from more standard usage.

### Interface Terms

**Combo Box**—An interface element which contains a dialog box and a pull-down menu as follows:

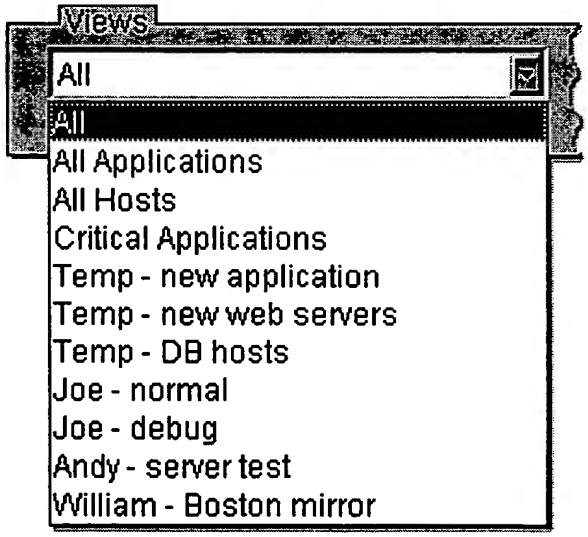


**Click**—Push the left mouse button over the specified object on the screen.

**Select**—Choose from a menu or set of options by clicking on the menu or set and dragging the mouse pointer down to the desired object.

### Web Site Terms

**Daemon**—A process that runs for a long period of time in the "background" (meaning that there is no direct control interface). While a daemon is running, it waits for specific conditions to occur and then performs a routine.

**Host**—The physical machine on which processes are run.

**HTTP**—*HyperText Transfer Protocol.* The method by which web content is passed over the internet.

**HTTP R direct**—A redirect is a specific response that a web server can give to a client. The response itself is the address of a different URL for the client to hit.

**IP Interface**—An interface on a host where a set of instructions (the *Internet Protocol*) allows processes on the host to interact with processes on other hosts.

**Port**—A specific location on an IP interface, represented by a number.

**Process**—An instance of a program running on a machine. In UNIX, processes can be listed by typing "ps" at the command line. In Windows NT, processes can be listed by calling the task manager (by pressing and holding <CTRL><ALT><DEL>).

**URI**—*Uniform Resource Identifier*, the part of a URL which specifies the path to information on the specified host.

**URL**—*Uniform Resource Locator*; an address widget that identifies a document or resource on the World Wide Web. A URL is a pointer either to a file on the local machine, or to a file elsewhere on the Internet.

**Web Server**—A daemon which waits for requests from a client and then returns files which contain HTML. In this guide, the terms "web server" and the more generic "server" are synonymous.

## Feedback

We want to hear from you! If you have questions or comments about the Web-Spective User's Guide, you can contact the Atreve Software Documentation Team in one of the following ways:


Fax: (617) 354-0513
E-Mail: documentation@atreve.com
Postal Mail:

Documentation Team c/o
Atreve Software, Inc.
767C Concord Avenue
Cambridge, MA. 02138


Note that there is also a survey supplied with this User's Guide that you can use to send comments to the Documentation Team.

# Chapter 1

## Product Overview

Atreve Software, Inc. proudly introduces its powerful new web management tool, WebSpective. Peak Web was designed to optimize sites that average ten thousand or more hits per day and that incorporate several machines for high availability. WebSpective turns the task of running multiple servers with varied content into an easy, single-user operation.

## Chapter Topics:

- What is WebSpective?
- Basic Concepts

## What is WebSpective?

WebSpective's overall operation is divided between three different focus areas:

### High Availability

By maintaining a continuous and dynamic watch over a web site's vital signs, WebSpective can deliver unprecedented availability and recovery. WebSpective directs clients away from failed servers and can attempt to restart web servers which have gone down. WebSpective's own design defends against software faults and can be configured with redundant components to further ensure crash resistance.

### Traffic Control

WebSpective's first line of hit management is called Interceptor. Interceptor runs both as part of WebSpective or as a stand-alone module. As a stand-alone, Interceptor takes incoming hits and distributes them across web servers using a simple load-balancing algorithm. As a stand-alone module, Interceptor can also be used to control general access to web content. Combined with the rest of the WebSpective package, Interceptor's range and availability are dramatically increased. Using dynamic performance information, WebSpective can continuously change hit distribution patterns to optimize the load across web servers.

### General Manageability

The third part of WebSpective's operation is a set of administration tools. The WebSpective architecture provides the user with a level of site management that is not available in today's market. WebSpective keeps a database of vital information about a site's servers and can be configured to take action in response to a number of hit- and system-related events. Comprehensive, dynamic server and content profiles are provided along with access to fundamental site tools in a single user interface.

## Basic Concepts

The Interceptor was designed to avoid the problems encountered by more traditional traffic management systems. To understand the Interceptor, you should first become familiar with the methods of *DNS Round-Robin* and the *Proxy Server.*

### DNS Round-Robin

DNS round-robin and load-balancing DNS are two mechanisms used to distributed incoming traffic over multiple machines. In a DNS round-robin setup, the DNS server is configured with several IP addresses that all correspond to the same host name. Each DNS client is passed the next IP address in the queue in a round-robin fashion.

Load balancing DNS is a slight improvement over round-robin DNS in that it provides a level of intelligence to traffic distribution by measuring machine load and sending the next request to the least busy system. A monitoring process is required to send machine usage information to the DNS.

Limitations:

- DNS-based solutions cannot detect when a Web server is down. Unfortunately, this situation can occur due to the inability of web servers to scale when overloaded. When this happens, the machine is up, but the Web server is down. DNS does not know when a machine is down, so it continues to use that machine despite the fact that it no longer can serve Web content. This results in bad connections for users.
- DNS round-robin has no knowledge of the machines whose IP addresses it serves. This means that all machines in a round-robin DNS record must have completely replicated content (as opposed to replicating only the most popular content). This is an expensive proposition for multi-server sites that have client connections to back-end RDBMS's and corporate applications. This configuration is also very inefficient and time consuming.

- DNS servers throughout the Internet cache IP addresses. If a Web server or machine goes down, these remote caches cannot detect this and will continue to send requests to the downed server, resulting in user connection failures. By the same token, if a new replicated Web server is added to the system, it can take hours and sometimes days for the new IP address to propagate around the Internet. Additionally, since IP addresses are cached by DNS, all users of a particular DNS server (as with users in the same corporation) get the same IP address. This results in the overload of a single machine.
- DNS poses an upper limit to scalability. DNS allows a maximum of 32 entries to be configured in the IP list. That means no more than 32 machines can be used under this scheme.

## Proxy Servers

The security, caching, and load balancing tasks at high-profile web sites has also been handled by proxy servers. A proxy server acts as a conduit for hit requests and redirects them internally to the appropriate server. Unfortunately, proxy servers can become a bottleneck, denying service to clients and slowing down under the burden of thousands of requests.

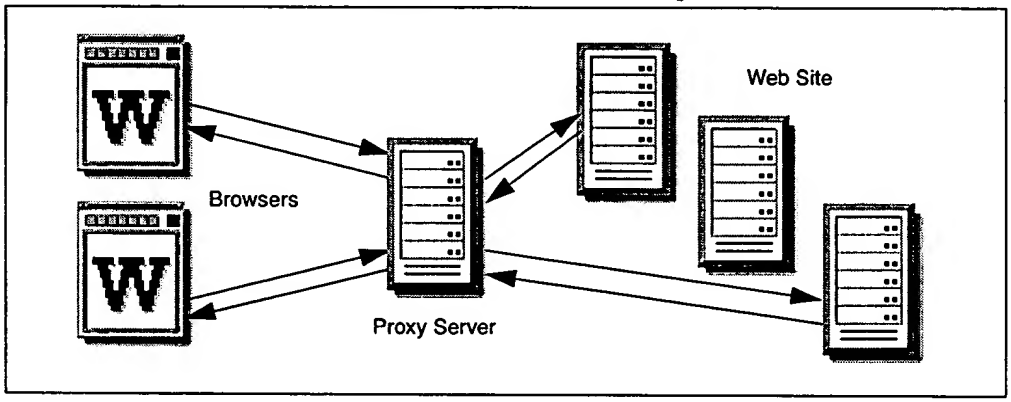


## The Interceptor

The Interceptor provides a new approach to traffic management that greatly reduces the occurrence of bottlenecks and lag time. By using HTTP redirec-

tion, the Interceptor spends much less time dealing with individual clients and does not have to establish connections directly to the web servers.

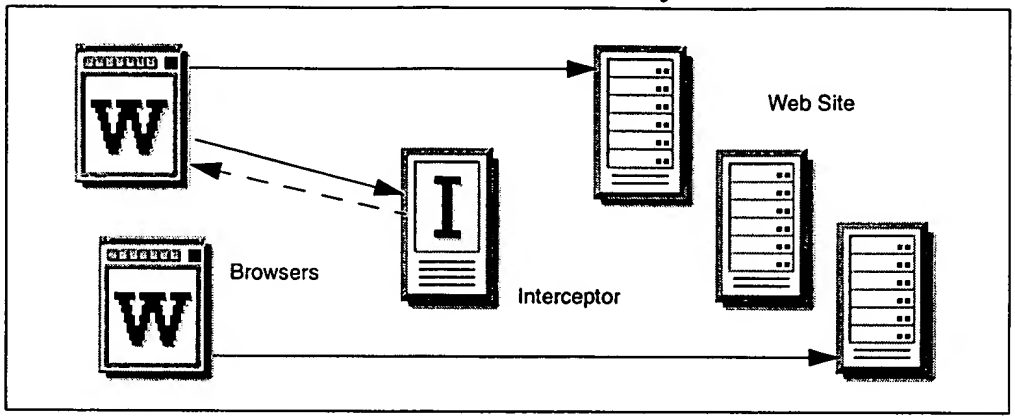


When a client hits the Interceptor, it refers to a list of hosts, webservers, and endpoints to determine the most likely recipient of the request. It then sends an HTTP redirect command to the client, which contacts the recommended server.

The Interceptor bases its redirection requests on information about the web site's servers and the machines they reside in. Initially, the system administrator provides this information. This static image of the web site is sufficient to make dramatic improvements to the site's hit response time.

With the release of WebSpective, the Interceptor's effectiveness will be greatly extended. WebSpective will continuously update Interceptor's view of the web site, allowing the software to make real-time decisions about how to best distribute hits that it receives.

# Chapter 2

## Installation and Configuration

## Chapter Topics:

- Installation Basics
- Windows NT Installation
- UNIX Installation
- Configuration

## Installation Basics

On UNIX systems, installation begins with the user running the pkgadd utility. On Windows NT systems, installation is managed by the Windows InstallShield.

### A Warning on Network Installations

It is important to note that all files should be installed locally. This means that files should not be installed on network drives, and should also not be installed on NFS mounted drives. Some files cannot be accessed properly if they are not on the local machine, and the practice of local installation further enhances overall site security.

### System Requirements

The Interceptor will run under Windows NT 4.0 (with Service Pack 3 or higher) and Solaris 2.5.1. While RAM requirements will vary between implementations, Atreve recommends 32Mb on Windows NT systems and 96Mb on UNIX systems.

## Windows NT Installation

On Windows NT systems, installation of the Interceptor is done with InstallShield. To begin installation, locate and run the Interceptor's "Setup.exe" file.

When you install the Interceptor, you will need to generate a keyfile (described in the section titled "The Keyfile" in Appendix B). The keyfile generation program requires you to enter a randomly timed string of characters as directed by the install tool.

In addition to putting files in their proper place, the installation tool sets configuration parameters that determine the behavior of the Interceptor during run time. The complete listing of configuration parameters can be found in Appendix A. Refer to the section entitled "Configuration" on page 8 for information on viewing and editing these configuration parameters before starting the Interceptor.

## UNIX Installation

Installation on UNIX (specifically Solaris) systems is done with the pkgadd utility. You can either run this directly from the WebSpective CD-ROM, or you can copy the necessary information to your local machine first.

### Local Installation

To perform the installation from your local machine, locate the file named "webspective1_0.tar.Z" on the media you received and copy it to a temporary directory. Uncompress and expand the file as follows:

```
uncompress webspective1_0.taz.Z <ENTER>
tar xvf webspective1_0.tar <ENTER>
```

The file expansion creates a new directory titled "/webspect". This directory contains all of the information that pkgadd will need to install the WebSpective system.

### Setting Up Files with *pkgadd*

Log in as root and go to the WebSpective package's base directory. If you are installing directly from the CD-ROM this is /solaris/webspect. Run pkgadd with the following command:

```
pkgadd -d . -R <load-path> <ENTER>
```

Where *load-path* is the path under which you would like to install WebSpective (usually /usr/local). The utility will warn you that it is making changes to the system as root. Specifically:

- The interceptor_exe and the Agent are setuid to root and setgid to group "nobody"
- The pkgadd tool will be running scripts as root

For information on disabling setuid and setgid, refer to Appendix D: Security Considerations.

Installation proceeds, and when it has finished, the path to a set of template files appears. These templates are the registry files for each component of the WebSpective system. Copy the template for the components you are installing to a permanent location. Atreve suggests the "etc" directory below the directory in which you installed WebSpective.

## Making a Keyfile

Before you can run your WebSpective component, you need to provide or generate a keyfile for the system (See Appendix B for a description of what the Keyfile is). Packaged with WebSpective is a utility called "keygen". Run the keygen utility by entering the following command:

```
keygen [keyfile.key]
```

Where *keyfile.key* is the name you want to assign to your keyfile. Put the keyfile in the desired location.

Once you have installed the components and prepared a keyfile, you are ready to edit the registry file of the component you wish to run. Refer to the section titled "Configuration" on page 8 for information on editing registry files.

## Setting Environment Variables

Before WebSpective will run properly, the following variables must be set in the user's shell.

- **LD_LIBRARY_PATH**—must be set to point to the location of the WebSpective shared linraries. In k-shell, this would be done as follows:

  ```
  export LD_LIBRARY_PATH="$LD_LIBRARY_PATH:/usr/local/atreve/
  lib"
  ```

- **NLSPATH**—must be set to point to the location of the WebSpective message catalog files. In k-shell, this would be done as follows:

  ```
  export NLSPATH="$NLSPATH:/usr/local/atreve/etc/%N.cat"
  ```

- **PATH**—must be set to point to the location of the WebSpective executables. In k-shell, this would be done as follows:

```
export PATH="$PATH:/usr/local/atreve/bin"
```

Once you have installed the Interceptor and prepared a KeyFile, you are ready to configure the registry file.

## Configuration

The Interceptor's registry settings determine many of the aspects of its behavior during run time. All of the registry parameters are listed and described in Appendix A.

On Windows NT systems, these settings are set during product installation and are stored in the program registry. If you are using an NT-based Interceptor, you do not need to do any additional configuration. Go on to the next chapter, "Starting and Running the Interceptor".

On UNIX systems, the registry settings must be set by the user. Provided with the Interceptor is a registry file template that you may need to edit to reflect your specific configuration.

### How to edit the registry template

The registry template is installed to the following location:

```
<load-path>/atreve/etc/templates/interceptor.reg
```

Where *load-path* is the path under which the interceptor was installed (/usr/local/ by default).

1. Copy the interceptor.reg file to your working area and open it. You will see that it contains a number of registry parameters set to default values.
2. Edit the necessary parameters. For instance, to set the Interceptor's host and port to www.xyz.com:5040, you would change the following parameters:

   ```
   LocalIP=www.xyz.com
   LocalPort=5040
   ```

   Appendix A specifies which parameters must be set correctly in order for the Interceptor to work.
3. Once you have finished editing the file, save it to the directory from which you will be starting the Interceptor.

When the registry file has been set correctly, you will be able to start the interceptor and begin adding applications for it to manage.

# Chapter 3

## Starting and Running the Interceptor

### Chapter Contents:



## Starting and Stopping the Interceptor

The Interceptor can be configured to run either on machine startup or at the explicit command of the user.

### Startup

On Windows NT systems, the Interceptor exists as a service that starts when the system is started. It can be controlled locally through the Services Control Panel.

**Local Startup.** On UNIX systems, the Interceptor must be started by the user unless otherwise directed. To start the interceptor, type the following at a command prompt on the Interceptor's local machine:

```
interceptor -r <filename.reg>
```

Where *filename.reg* is the name of the file that contains the Interceptor's configuration parameters.

**Boot Startup.** If you would prefer to configure your UNIX-based Interceptor to start at system startup, we recommend that you place the script/commands within "/etc/rc3.d" after S25. You could set up "/etc/rc3.d/S25atreve", for example, but check the file "/etc/rc3.d/README" for clarification. At this point your system will have "syslogd" running and the file systems will be exported.

Note that all shared libraries used by our process (under ../atreve/lib) need to be symbolically linked from within "/usr/lib" as the set user ID option only

looks for shared libraries under "/usr/lib". Also note that any process (either a script or program) started through "/etc/rc*" will run as root unless you explicitly use "su" to specify the process id

**Checking the Startup.** You can check to see that the Interceptor has started successfully by checking the UNIX syslog or the Windows NT Event Viewer. both the Interceptor and its watcher program log messages about the startup.

### Shutdown

The Interceptor can be shut down in a few different ways. First, you can stop the Interceptor through the Interceptor Control Program (ICP). Next, on UNIX systems, the Interceptor can be stopped with a plain "kill" command either on its watcher process or on the Interceptor itself. Finally, on Windows NT systems, the Interceptor can be stopped from the Services Control Panel.

We recommend attempting to stop the Interceptor via the ICP before shutting it down locally.

## The Interceptor Control Program (ICP)

In the WebSpective system, Interceptor re-configuration is handled automatically by the Manager. As a stand-alone product, the Interceptor is controlled by via the Interceptor Control Program.

#### How to start the Interceptor control driver

To start the ICP and begin interacting with the Interceptor, enter the following:

```
interceptor_cp -k keyfile [-c cmd_file]
```

In this command, *keyfile* is the path, in either relative or absolute terms, to the keyfile, and *cmd_file* is a multiple-command file for the ICP to process. For more information on the multiple-command file, refer to the section titled "Executing Command Files".

Note that in order for the control driver to work properly, the shell in which it is executed must be able to access the keyfile.

### Connecting to the Interceptor

These commands are system-oriented commands for setting up and exiting the ICP.

- **help, ?**—Print the command list

Note: Including the name of a command in the help query will yield information on that specific command.

- **pen**—Establish a connection to the Interceptor
- **status**—prints the current Interceptor host and port configuration
- **quit, exit, bye**—Quit the Interceptor control driver.

### How to set the Interceptor's host and port

Before you can work with the Interceptor, you must be able to establish a connection with it.

1. At the driver prompt, type "open" and press <ENTER>.
2. The ICP asks for the name of the host on which the Interceptor resides. Type the name of the Interceptor's host and press <ENTER>.
3. The ICP asks for the port number which the Interceptor has been assigned. Enter the port number and press <ENTER>.
4. The ICP returns you to the command prompt. You are now ready to send commands to the Interceptor. To verify your host and port settings, type "status" at the prompt and press <ENTER>. To verify your connection to the Interceptor, use the "query" command.

## Interceptor Run-Time Control

This section contains procedures for running the tasks which are available to the user through the command driver. See the "The Interceptor Control Program (ICP)" section on 10 for information on starting the ICP.

### Command Listing

These commands are used to control the Interceptor through the ICP.

- **add app**—Add a new application
- **add endpoint**—Add a new endpoint
- **delete app**—Delete an application
- **delete endpoint**—Delete a endpoint
- **deactivate**—Shut down an application or enitre web site
- **activate**—Restart an application or enitre web site
- **set threads**—Change the number of threads given to an application
- **set strength**—Assign a new static strength to a server
- **query**—Retrieve a list of current applications and servers
- **save**—Write changes to the Interceptor's configuration file

All of the commands can be entered either one element at a time or with their complete set of arguments on a single command line. Note, however, that if

arguments are entered in the wrong order that the control driver will not understand the command.

Note that new values for Interceptor properties must follw these rules:

- Names cannot contain spaces
- Names cannot exceed 64 characters in length
- Characters used must be from the ASCII standard character set

### How to view the Interceptor's current configuration

Before and during a re-configuration process, you may want to see the Interceptor's configuration. By typing "query" at any time at the ICP prompt, you can get a complete listing that includes the Interceptor's system configuration and information on all of the applications and endpoints in the Interceptor's load-balancing routines.

### How to add an application to the Interceptor's register

The Interceptor classifies a web site according to different kinds of content. An example would be an Internet Service Provider (ISP) that supports a number of different companies on a combination of hosts. Each company is considered an *application*, and all of the servers managed by the Interceptor are categorized by the application they carry. Consequently, applications must be defined before servers can be added to them.

*Single Entry Listing*

To submit this command on a single command line, type the following and press <ENTER>:

```
add app app-name http-host:http-port threads secure [SSL-port]
```

*Step-by-Step Operation*

1. At the ICP prompt, type "add app" and press <ENTER>.
2. Enter a name for the new application and press <ENTER>. The name you choose is for your own reference, and can be up to 256 characters. Note that the name can not contain the forward-slash (/) symbol as this is reserved for use by the Interceptor.
3. Enter the host machine's name (as in "host.domain.com") and press <ENTER>.
4. Type the port number at which the application is located on the host machine and press <ENTER>.
5. The ICP asks you for the number of threads to assign to the application. The number of threads an application has is equivalent to the number of clients who can be simultaneously redirected by the Interceptor. Consequently, the number that you assign to an application is highly dependent on the number of hits that the application is expected to receive. The maximum number that you can assign to any given application is governed by the operating system on which the application is running. Type a thread count and press <ENTER>.

6. Indicate whether the application will be running on a secure server by typing "1" for secure and "0" for insecure and pressing <ENTER>.
7. If you are setting up a secure application, you may then enter the secure server's port number. If you do not provide a value for the secure port number, the Interceptor assumes that the initial SSL connection is itself insecure. Type a port number or leave the value blank and press <ENTER>.
8. The ICP returns you to the driver prompt. The application is now registered with the Interceptor. You can confirm this by typing "query" and pressing <ENTER>.

### How to remove an application from the Interceptor's register

Removing an application completely erases all of the information regarding the application and its servers from the Interceptor's register. If you only want to temporarily disable the application, see "How to deactivate an application or an entire web site".

*Single Entry Listing*

To submit this command on a single command line, type the following and press <ENTER>:

```
del app app-name
```

*Step-by-Step Operation*

1. Type "del app" at the ICP prompt and press <ENTER>.
2. Type the name of the application that you want to remove and press <ENTER>. The command driver removes the application from the Interceptor's register.

### How to add an endpoint to an application

Once an application has been defined, follow these steps to define the endpoints which will be associated with it.

*Single Entry Listing*

To submit this command on a single command line, type the following and press <ENTER>:

```
add endpoint app-name endpoint-host:endpoint-port strength
```

*Step-by-Step Operation*

1. At the ICP prompt, type "add endpoint" and press <ENTER>.
2. Type the name of the application to which you wish to add an endpoint and press <ENTER>.
3. The ICP prompts you for the name or IP address of the endpoint's host. Type either the full host name or IP address and press <ENTER>. Note: If the AccDomains parameter is set, you must use the host name, and not the IP address.
4. Type the port number at which the endpoint is located on the host machine and press <ENTER>.

5. The ICP prompts you to enter the relative strength of the server. The range is 1 to 10. The value is determined by measuring the server's strength against other servers in your web site. One rule of thumb is to base the number on how many requests a server can process in a given unit of time. Type in a relative strength and press <ENTER>.
6. The ICP returns you to the driver prompt. The server is now registered under the chosen application.

Note: In a multi-IP (multihoming) environment, a single server "listens" to several ports. The Interceptor sees each endpoint (server host and port) as a different entity.

### How to remove an endpoint from an application

Removing an endpoint from an application completely erases that endpoint from the Interceptor's register. The removal of an endpoint additionally causes the server which supports the endpoint to be removed from its host. Note that once an endpoint has been removed, clients who have bookmarked that endpoint will not be redirected by the Interceptor.

*Single Entry Listing*

To submit this command on a single command line, type the following and press <ENTER>:

```
del endpoint application-name endpoint-host:endpoint-port
```

*Step-by-Step Operation*

1. Type "del endpoint" at the ICP prompt and press <ENTER>.
2. Type the name of the application which the endpoint is associated with and press <ENTER>.
3. The ICP prompts you for the name or IP address of the endpoint's host. Type a name or IP address and press <ENTER>.
4. Type the number of the port at which the endpoint is located on the host machine and press <ENTER>.
5. The ICP removes the endpoint from the application listing and returns you to the driver prompt.

### How to change the thread count for an application

If the thread count for an application is too high, then the server wastes resources which might be better given to other applications. If the thread count is too low, the number of service-denied clients will be high. You can adjust the thread count for a given application to compensate in either case.

*Single Entry Listing*

To submit this command on a single command line, type the following and press <ENTER>:

```
set threads app-name new-thread-count
```

*Step-by-Step Operation*

1. Type "set threads" at the ICP prompt and press <ENTER>.
2. Type the name of the application for which you want to change the thread count and press <ENTER>.
3. Type the new thread count and press <ENTER>.

### How to deactivate an application or an entire web site

When you deactivate an application or the web site, the control driver will prompt you for a URL (either a web page on a remote server or, more likely, a local file) which contains "sorry page" content. The Interceptor will guide any requests for the shut down application or site to this URL.

If you have previously entered a shutdown URL for a given application, you do not need to re-specify it the next time you shut the application down. Leave the shutdown URL entry blank when issuing the Deactivate command and the ICP will use the established shutdown URL.

*Single Entry Listing*

To submit this command on a single command line, type the following and press <ENTER>:

```
shutdown app-name (or "null_app") shutdown-URL
```

*Step-by-Step Operation*

1. At the ICP prompt, type "deactivate" and press <ENTER>.
2. Enter the name of the application you wish to shut down. If you intend to shut down the entire web site, enter "null_app" as your application name. Press <ENTER>.
3. The ICP prompts you for a URL which the Interceptor will use to redirect clients. Enter the URL and press <ENTER>.

### How to reactivate an application or web site

When you are ready to reopen an application or the entire web site, do the following:

*Single Entry Listing*

To submit this command on a single command line, type the following and press <ENTER>:

```
activate app-name (or "null_app")
```

*Step-by-Step Operation*

1. Type "activate" at the ICP prompt and press <ENTER>.
2. Enter the name of the application to restart, or type "null_app" to restart the web site. Press <ENTER>.

### How to save changes t the configuration file

Changes you have made to the Interceptor's configuration file are lost once the Interceptor is shut down. To avoid losing these changes, you can save them to the configuration file. You should back up your old configuration file as a safety precaution before you do this.

1. Type "save" at the ICP prompt and press <ENTER>.
2. Changes you have made to the config file are automatically written to disk.

## Executing Command Files

A command file is a simple of way of performing a number of tasks on ICP start-up. The command file can contain any number of commands, but must begin by opening a connection with the Interceptor, and must end by saving the changes to the Interceptor. Once the ICP is done processing it automatically exits. Consequently, if you want to continue to use the ICP you will need to restart it without loading a command file.

Here is an example of a command file:

```
open hostname:port
query
add app test_app01 sixpack.atreve.com:18701 12
add endpoint test_app01 app01a.atreve.com:81 5
add endpoint test_app01 app01b.atreve.com:81 5
add endpoint test_app01 app01c.atreve.com:81 5
add endpoint test_app01 app01d.atreve.com:81 5
save
exit
```

Note that comments are not currently accepted but blank lines are OK.

# System Behavior

## Logged Events

On Windows NT systems, the following events are recorded to the Event Log. On UNIX systems, they are recorded to the syslog.

- Interceptor start
- Interceptor shutdown
- Authentication failure
- Timeout
- Adding or deleting an application
- Adding or deleting an endpoint
- Changing parameters (any of them, at any level)

## Client Connections

If a request arrives on the endpoint for an application, the Interceptor processes it according to these rules:

- If the application is Enabled, and there is at least one server for it, then the Interceptor replies with an HTTP "302" redirection request to one of those servers. It chooses the server based on the Interceptor's load-balancing algorithms (see Appendix B for more information).
- If the application is Enabled, but there are no servers for it, then either the contents of the "Sorry Page" are returned as content. If that parameter is not defined, a generic "503" error message saying that no servers are available will be returned.
- If the application is Disabled, then either the contents of the "Sorry Page" are returned, or, if that is not defined, a generic "500" error message saying that the site has been temporarily disabled will be returned.

## Failure Tolerance

If the Interceptor fails, it will attempt to restart itself. If the attempt to restart fails, then requests to the Interceptor cannot be served. However, the application servers will remain active, and old connections and bookmarks to these servers will still be functional.

# APPENDIXES

# Appendix A: Registry Parameters

This Appendix contains a listing of all of the registry parameters used by the Interceptor. Some parameters only affect the Interceptor when used as part of the WebSpective system. These parameters have a (WEBSPECTIVE) tag.

**AccDomains**

*Optional*: Key (a text string) and DNS name

When set, the Interceptor will not accept the introduction of servers outside of the given domains. The AccDomains parameter is a section, inside of which are keys and DNS names. The key "Atreve Software, Inc.," for instance, might have the value "atreve.com," in which case servers hosted at ws3.atreve.com would be permitted service. If a value is not set for this parameter, you will be able to add any servers under any domain name to the system. For example:

```
XYZ Company=xyz.com
```

This entry will allow any machine under xyz.com to be registered with the Interceptor.

```
WXY Systems=web.wxy.com
```

This entry will only allow servers running on web.wxy.com to register with the Interceptor.

**AcceptSSL**

*Optional*: True (1) or False (0)

*Default*: 0

This parameter specifies whether or not the Interceptor will accept SSL requests. By default, it does not. Note that other SSL related parameters will not work unless this is made true.

**AuditCommands**

*Optional*: True (1) or False (0)

*Default*: 0

This parameter controls the level of system logging. If this parameter is made true, then all commands received by the Interceptor are logged with the client IP address, time/date stamp, and a message to the system log. if this parameter is made false, only Interceptor startup, shutdown, and error messages are logged.

**CertificateFile**

*Optional*: Absolute path to SSL certificate

Parameter containing the path to the file containing the SSL certificate for your site. Note that this must be the same certificate that your servers use, in order for the Interceptor to pass requests to them seamlessly. Also note that the AcceptSSL parameter must be set to true for the Interceptor to handle SSL requests.

**Cl seD lay**

*Optional*: Integer

*Default*: 0 on Windows NT systems, 1 on UNIX systems

This optional parameter specifies a time to wait, in seconds, before closing a connection at the end of a session. On some TCP/IP implementations, allowing the client to close first frees resources faster.

**CloseTimeout**

*Optional*: Integer

*Default*: 5

This optional parameter specifies the number of seconds which the Interceptor will wait before closing a connection to a client, after writing data to that client. Note the difference between this parameter and RecvTimeout, which is an equivalent waiting period to read data from the connection

**DecayWindow**

*Value*: Integer

*Default*: 600

*Used by*: Interceptor

(WEBSPECTIVE) The amount of time in seconds that the Interceptor will consider any one set of load data accurate. If the Manager does not send an update before the DecayWindow expires, the Interceptor will return to the default settings that it loaded on start-up.

**Description**

*Optional*: Text string

This optional parameter is an arbitrary string for users that describes this Interceptor instance

**KeyFile**

*Required*: Absolute path to keyfile

This is a required parameter. This parameter's value is the name of a file to use as a secret key. The key itself is generated during the installation and must be identical for each of the various WebSpective components. The keyfile is not used in processing HTTP requests.

**LocalIP**

*Required*: Interceptor's IP address or DNS listed name

This parameter specifies the IP interface at which the Interceptor listens for connections from the ICP (or the Manager).

**LocalPort**

*Optional*: Port number

*Default*: 4040

The IP port on which the Interceptor listens for connections from the interceptor_cp (or the Manager). If this parameter is not set, then the Interceptor will perform as well as the static system information permits.

**Log**

*Optional*: Path to desired log file

When specified, the Log parameter provides the Interceptor with the name of a file to which it will log events. If this parameter is not specified, the

interceptor will default to the system event log—"syslog" on UNIX systems and the Event Viewer on Windows NT systems.

**ManagerIP**

*Required*: Manager's IP address or DNS listed name

(WEBSPECTIVE) This parameter specifies the IP interface at which the Manager listens for connections from other WebSpective components.

**ManagerPort**

*Optional*: Port number

*Default*: 4040

(WEBSPECTIVE) The IP port on which the Manager listens for connections from other WebSpective components.

**ReadOnly**

*Optional*: True (1) or False (0)

*Default*: 0

A true/false parameter. If this parameter is set, the Interceptor will not update its configuration as it runs. This provides better data security to the configuration, but may cause sub-optimal system performance.

(WEBSPECTIVE) If the Interceptor's changes at run-time are lost (due to the Interceptor's machine crashing, for instance), then a restarted Interceptor will not receive dynamic configuration changes until a connection is made to the Manager.

**RecvTimeout**

*Optional*: Integer

*Default*: 5

This is an optional parameter which specifies how long to wait, in seconds, for data from a client before giving up and disconnecting. Too low a value will result in clients on slow network connections having difficulty establishing a connection; too high a value may result in many threads waiting for user traffic, and will create a security vulnerability to a denial-of-service attack.

**ReusePort**

*Optional*: True (1) or False (0)

*Default*: 0

This optional parameter controls whether or not the Interceptor will try to use a listening port if it is already in use. If this parameter is set to False (0), the Interceptor will not start if the port is in use. If the parameter is set to True (1), the Interceptor will start up expecting the current processes to die and leave the Interceptor in place.

**SecureMessageTimeout**

*Optional*: Integer

*Default*: 60

(WEBSPECTIVE) The age for which a message to the Interceptor should be accepted from other components in the WebSpective system. This parameter has been established as a defense against "replay" attacks. The value is a number of seconds. When entering a value, consider possible differences in system clock readings and network latency. Otherwise you run the risk

of invalidating all of the messages that other components send to the Interceptor.

**SSLPassword**

*Optional*: Text string

This parameter contains the Interceptor's secure server password (if an application requires it). Note that the AcceptSSL parameter must be set to true for the Interceptor to handle SSL requests.

**UseRefreshPage**

*Optional*: True (1) or False (0)

This setting applies only to requests recieved on the SSL port of a secure application. If an interceptor does not have any secure applications using SSL, then this setting can be ignored.

For applications using SSL this setting controls whether or not a page is sent to the browser when redirecting the SSL request to a server. Older browsers which do not support automatic redirection require a page to be sent with a link for the user to follow into the SSL content.

Most browsers now support automatically following the redirected request, however, if the site wishes to support older browsers then this option should be set to 1. The following browsers do not need this option to be set: Microsoft Internet Explorer 3.0 and later, Netscape 2.0 and later.

# Appendix B: Security Considerations

This appendix is intended to provide you with a basic discussion of WebSpective's security infrastructure. Additionally, you will find information on how to make WebSpective conform to your organization's security policies and standards.

## The Keyfile

During installation, you will be asked to generate or locate a *keyfile*. The keyfile contains a random string of code that the Interceptor uses to verify commands sent by other components. In this fashion, intruders cannot mimic the WebSpective system and send faulty commands or information.

The keyfile needs to be generated only *once*. Once created, access to the keyfile should be controlled so that only WebSpective components can read it, and nothing can write to it.

## Component Communication

The HTTP ports to which the Interceptor listens are protected against hung or invalid requests and denial-of service attacks (see the CloseDelay property on page 2). Connections between components are authenticated and validated (to prevent spooling attacks) with the KeyFile. None of these connections are encrypted.

## Components and *setuid* (UNIX systems only)

Often, WebSpective components will be configured to listen on the default HTTP and HTTPS ports, 80 and 443. To acquire these resources, the components must have administrative permissions. This is accomplished through the use of *setuid* and *setgid*, which allow the components to acquire the ports a root and then demote themselves to user "nobody".

If you do not want the Interceptor and Agents to use setuid and setgid during their operation, then you have two options. You must either run these process as root or assign them to ports which do not require root permission. Additionally, you must ensure that your web servers are in the same group as the Agent—otherwise, the Agent will not be able to read Filter information from shared memory.

## Filesystem security

WebSpective configuration files may include sensitive information, such as database passwords and secret keys for authentication. Similarly, the running memory images of the components may have sensitive data within them. It is the site administrator's job to maintain physical security of the host and of its filesystem to be sure these resources are not compromised.

# Appendix C: The FAQ

1. What UNIX environment variables need to be set to run the Interceptor?

   An assumption is that the product is installed under the directory "/usr/local/atreve". The following variables must be set. With k-shell the commands are:

   ```
   export LD_LIBRARY_PATH=/usr/local/atreve/lib:$LD_LIBRARY_PATH
   export PATH=/usr/local/atreve/bin:$PATH
   export NLSPATH=/usr/local/atreve/etc/%N.cat:$NLSPATH
   ```

2. What is the Interceptor Watcher program?

   The Interceptor Watcher program is the parent process that monitors and restarts the Interceptor program as necessary.

3. What is the Interceptor Control Program (ICP)?

   The ICP is a command line driver that enables the user to interactively configure an application on the Interceptor. See the product documentation for details and command lists.

   Note that a valid security key file must be supplied to communicate with the Interceptor. I.e. both must be started with the same key file. Copies of a single security key file can be made and stored separately.

4. What modifications need to be made to run the Interceptor and Interceptor Watcher programs as root?

   The set-user-ID ("setuid") bits for the Interceptor executable ("interceptor_exe") and the Interceptor Watcher program ("Interceptor1_0Interceptor") need to be set and the files owned by user root.

   Note: as the set-user-ID bit will only accept shared libraries from "/usr/lib" symbolically, links need to be made from "/usr/lib" to "/usr/local/atreve/lib" for all the libraries within "/usr/local/atreve/lib".

   NOTE that the next Interceptor version provides this option as a part of the installation process.

5. How do I start the Interceptor on Solaris?

   Follow the installation process described in the documentation. An assumption is that the product is installed under the directory "/usr/local/atreve".

Copy the template registry files from "/usr/local/atreve/etc/templates" to your working area. For the Interceptor there will be a sample file named "interceptor.reg"

Now use keygen to generate a security key.

```
keygen  test.key
```

Starting the Interceptor:

Edit the "interceptor.reg" and modify the following options as required:

```
LocalIP=<hostname> // if you don't want to use the default
LocalPort=<default> // if you don't want to use the default
Keyfile// specify the location and name of the keyfile
[WebSpective Interceptor/1.0/Interceptor/AccDomains]
// if you do not specify a domain by default all domains are
accepted.
// domain specification is done as shown below.  This will only
allow endpoints
// within the atreve domain to be specified.
atreve.com=atreve.com


[WebSpective Interceptor/1.0/Interceptor/Startup]
// no change required unless the executable "interceptor_exe"
is not in the following path.
CommandLine = /usr/local/atreve/bin/interceptor_exe


[WebSpective Interceptor/1.0/Interceptor/Managers/Manager]
// no change required in standalone mode
ManagerIP=<hostname>
ManagerPort=<default>
```

Now start the Interceptor as follows:

```
interceptor -r interceptor.reg
```

6. How do I know that the Interceptor started successfully?

   Both the Watcher program and the Interceptor programs log messages to the UNIX system log as daemon processes. I.e. via "syslogd". So a "tail" of the "syslog" should show the appropriate messages.

   Note that the next version of the Interceptor writes start up information to stdout.

7. How do I automate the Interceptor start process on Solaris?

It is recommended that you create a script that

8. How do I automate the Interceptor start process on NT?

   The WebSpective installation process installs the Interceptor as a NT service with the appropriate entries in the NT registry.

9. Where does the Interceptor and Watcher program log information?

   Within the UNIX system log. Both the Watcher program and the Interceptor programs log messages to the UNIX system log as daemon processes. I.e. via "syslogd". So a "tail" of the "syslog" should show the appropriate messages.

   Verify that the file "/etc/syslog.conf" is configured correctly to accept daemon process messages. The default Solaris system settings accept daemon messages.

   Note that you can modify the "/etc/syslog.conf" as follows to send messages to a specific log file:

   ```
   daemon.err;daemon.warning;daemon.info;auth.notice
   /var/adm/atreve.log
   ```

   Note a tab character is a REQUIRED separator between the ".;auth.notice" and "/var/adm/atreve.log" text.

   Note that for the "syslogd" process to accept changes in "/etc/syslog.conf" it must be sent a SIGHUP.

10. How do I stop the Interceptor program?

    The Interceptor can be stopped in the following ways. I.e. this will shutdown both the Interceptor and the Watcher programs. Issue a plain "kill" signal to the Interceptor process. I.e. "kill <interceptor pid>". Issue a plain "kill" signal to the Interceptor Watcher process. I.e. "kill <watcher pid>". Use the "shutdown" command from within the Interceptor Control Program (ICP).

11. How do I re-start the Interceptor program?

    The Interceptor can be re-started in the following ways.

    - Issue a "kill -1" (SIGHUP) to the Watcher program.
    - Issue a "kill -1" (SIGHUP) to the Interceptor program.
    - Issue a "kill -3" (SIGHUP) to the Interceptor program.

    We recommend that you use the "kill -1" to the Interceptor program.

    Note that the next version of the Interceptor has the "restart" option in the Interceptor Control Program.

12. Suggestions on administering the Interceptor?

    If both the Watcher and Interceptor programs run as root a "restart" has to be issued as a SIGNAL by a root process. I.e. either via a root account or a

"sudo" command. The other choice is to run the Watcher as a non-root process thus enabling administration of restart without having access to root privileges.

Note the next version of the Interceptor has the restart option available within the Interceptor Control Program thus providing for remote "restart" or "shutdown".

13. Can I use a command file with the Interceptor Control Program?

Yes. Use the following command:

```
interceptor_cp -k <keyfile> -c <command program>open six-
pack:18041
```

An example of the command program structure is as follows:

```
open hostname:port
query
add app test_app01 sixpack.atreve.com:18701 12
add endpoint test_app01 app01a.atreve.com:81 5
add endpoint test_app01 app01b.atreve.com:81 5
add endpoint test_app01 app01c.atreve.com:81 5
add endpoint test_app01 app01d.atreve.com:81 5
save
exit
```

Note that comments are not currently accepted but blank lines are OK.

14. Can you load a command file when you are within the Interceptor Control Program?

No. This is not supported at present.

15. Can you save to a separate configuration file?

No. This is not supported at present.

16. How do I change the strength of an Endpoint?

Use the "set strength" command. Use the "update" command. Refer to the Interceptor documentation for more details on these commands.

17. Can configuration file changes be undone?

I.e. if you delete an application can it be undone? Any changes since the last save command can be undone by restarting the Interceptor.

18. Can wildcards be used when defining applications or endpoints?

No. This is not supported at present.

# Glossary

**Application:** Essentially, a tree of web server content, such as might be served by a simple server. For fault tolerance and load-balancing, one might want multiple servers providing the same application. In the presence of multi-homing, one server might provide several applications.

A set of equivalent endpoints, each having the same content. Multi-homing machines provide edpoints for more than one application.

**Endpoint:** The terminus of a virtual circuit; a specific IP address and port pair. This might be used to describe the location at which a server is listening (with no circuit established), or one end of an established connection. In the second case, the two endpoints of the connection define which circuit is under discussion.

**Hardware virtual servers:** A form of multi-homing by which a single server process may imitate multiple servers. The server listens simultaneously on several IP interfaces, and the interface of the connection affects the server's behavior. Thus, for NetScape servers (from which the term is borrowed), the content offered varies based on the interface.

**Multi-homing:** As a generic, this can be any method of offering the effect of multiple servers on a single machine. This might be by hardware virtual servers, or by software virtual servers, or it might be simply by having multiple server processes on the same host.

**Software virtual servers:** A form of multi-homing in which a single NetScape server process, with a single listening endpoint, changes its behavior based on the host name in the URL on which the connection was made. This is impossible to detect from the server's perspective, unless the browser makes use of the optional field--and relies on the use of the optional "Host" header field in the HTTP request.

Multi-Process

Multi-homing in which independant web server processes each listen to a single endpoint. The servers provide unrelated content and are unaware of the other servers.

# Index

# Documentation Survey

Please help us to make our documentation better. If you have come across anything that you particularly like or dislike about the documentation for this product, please take the time to fill out this form and mail it to:

Documentation Dept. c/o
Atreve Software, Inc.
767C Concord Avenue
Cambridge, MA 02138

Thank you for you time.

---

Name: _____

Position: _____

Company: _____

What did you like about this documentation?

What didn't you like about this documentation?